Practitioner Brief

AI IN ASSET MANAGEMENT

Machine Learning in Commodity Futures: Bridging Data, Theory, and Return Predictability

Tony Guida

Practitioner Brief written by Cathy Scott

CFA Institute Research Foundation

CFA Institute Research & Policy Center

Machine learning (ML) has transformed equity investing, powering advances in factor discovery, portfolio construction, and systematic strategies. But commodities have not kept pace. Although commodities are widely traded and play a growing role in institutional portfolios, ML applied to commodity futures remains an underexplored frontier.

Commodity futures present different challenges than equities, but they also provide fertile ground for ML to add value. This chapter of *AI in Asset Management: Tools, Applications, and Frontiers* illustrates how ML, when grounded in commodity theory, can uncover persistent return patterns and convert the complexities of commodity markets into systematic, interpretable, and investable strategies.

By building features grounded in economics, applying supervised learning to rank commodities cross-sectionally, and combining models across horizons with ensembles, practitioners can create systematic portfolios that are both robust and interpretable.

Who Should Read This Chapter?

- Portfolio managers and chief investment officers (CIOs) seeking systematic commodity strategies that extend beyond passive diversification and inflation hedging.
- Risk officers, including chief risk officers (CROs), focused on horizon diversification, fragility reduction, and drawdown control through ensemble modeling.
- Quant researchers and data scientists interested in feature engineering tied to commodity economics and model interpretability.
- Asset allocators evaluating commodities as active sources of return, not only as defensive diversifiers.

Why This Chapter Matters Now

As ML continues to transform finance, the opportunities to apply them for commodities continue to grow.

- **Commodities have financialized.** Institutional flows, index-linked products, and hedge fund participation make commodities part of mainstream asset allocation.
- **Macroeconomic volatility is back.** Inflation shocks, energy crises, and supply chain disruptions have put commodities at the center of strategic conversations.
- **Machine learning has matured.** Equity investors already rely on ML pipelines, and extending these to commodities leverages proven methods in a new domain.
- **AI adoption is accelerating.** Investment firms are rapidly scaling machine learning, and commodities remain one of the last frontiers where these methods have yet to be applied.

"As systematic investing continues to evolve, the fusion of domain theory with modern ML holds the key to unlocking latent structure in every markets."

Tony Guida

Practical Applications

- **Build cross-sectional long–short portfolios:** Instead of forecasting absolute prices, use ML to rank commodities relative to one another. Models trained on features like momentum, basis, carry, and skewness generate scores that sort contracts from most to least attractive. Long positions in top-ranked commodities and shorts in bottom-ranked ones create systematic baskets that translate signals into investable strategies. This approach is especially valuable when commodities diverge from equities and bonds during macro dislocations.
- **Deploy ensemble models across horizons:** Short-term models may collapse in volatile markets, whereas long-term ones can lag in fast shifts. Ensembles solve this by blending predictions across horizons into a single score. This diversification across timescales smooths returns, lowers volatility, and reduces drawdowns. Ensemble portfolios provide more resilient risk–return profiles and greater scalability for institutional use.
- **Engineer features from commodity economics:** Models work best when features are grounded in theory. Commodity economics supplies a strong toolkit: basis from the Theory of Storage, carry from hedging pressure, skewness as a proxy for tail risk, and open interest as a measure of liquidity. Designing inputs around these signals makes outputs interpretable and easier to explain to boards, regulators, and clients.
- **Monitor macro-linked portfolio tilts:** Machine-learning portfolios often reveal exposures that mirror macro cycles. Energy contracts tilt long during geopolitical stress, grains during supply shortages, and metals in inflationary regimes. Monitoring these tilts gives CIOs and strategists early warning of regime shifts and offers a complementary risk overlay for broader allocations.
- **Apply liquidity and turnover controls:** Strong signals are worthless if they cannot be traded. Thin contracts like orange juice or rubber pose execution challenges at scale. Integrating liquidity screens and turnover constraints into portfolio construction ensures that ML strategies remain scalable and cost-efficient. Explicit modeling of slippage and transaction costs helps preserve alpha.
- **Detect regime shifts with feature monitoring:** Feature importance and SHAP values can double as early-warning tools. Rising skewness importance may flag tail risk, while shifts in basis or carry may signal inventory stress. Monitoring these changes in real time helps identify regime transitions before they show up in prices, giving risk teams and strategists a forward-looking edge.

Practitioner Toolkit

Role	Tool/Technique	Application	Benefit
Portfolio managers and CIOs	Cross-sectional ranking dashboards; long-short portfolio simulators	Rank commodities by ML scores and construct portfolios	Actionable, interpretable signals for allocation
Risk managers/ CROs	Horizon correlation monitors; ensemble risk analytics	Track diversification across horizons; monitor volatility and drawdowns	Reduce fragility; improve drawdown control
Quant researchers/ data scientists	Feature libraries (momentum, basis, skewness, open interest); SHAP pipelines	Engineer theory-based features; attribute predictions	Transparent, explainable ML models grounded in economics
Traders/execution teams	Turnover and slippage calculators; liquidity screens	Monitor trading costs; assess scalability in thin contracts	Ensure strategies are tradable at institutional AUM
Strategists/CIOs	Factor alignment reporting; macro-scenario stress tests	Validate ML signals against commodity factors and macro regimes	Governance-ready evidence of economic consistency

Workflow

1. **Define.** Select the commodity universe across energy, metals, and agriculture.
2. **Engineer.** Build features grounded in commodity theory, including momentum, basis, carry, skewness, and open interest.
3. **Train.** Apply supervised machine-learning models at multiple horizons, producing cross-sectional rankings.
4. **Aggregate.** Combine outputs across horizons into ensemble predictions to reduce fragility.
5. **Implement.** Translate rankings into long-short portfolios, applying liquidity screens and turnover controls.
6. **Monitor.** Track IC, SHAP values, turnover, and horizon correlations; validate signals against known factors and macro regimes.

The Bottom Line

Machine learning has been slower to take hold in commodity futures than in equities. For practitioners, that gap is less a limitation than an opportunity. The structural forces that once made commodities difficult for ML (inventories, storage costs, supply shocks, and hedging pressure) can now be converted into signals when models are grounded in economic theory.

For investment professionals, three lessons stand out:

- **Ground models in commodity economics.** Features such as momentum, basis, carry, and skewness are not ad hoc; they reflect real market structures.
- **Diversify across horizons.** Ensembles mitigate fragility and balance short-term alpha with long-term resilience.
- **Prioritize interpretability.** Linking ML outputs to established theories ensures strategies are explainable and defensible.

Commodities remain essential as diversifiers and inflation hedges. What ML adds is the ability to treat them not only as defensive exposures but also as systematic alpha generators. For portfolio managers, CIOs, risk officers, and quants, this represents a frontier of opportunity ready to be deployed.

Metrics That Matter

Information coefficient (IC). Measures the correlation between predicted rankings and realized returns. A positive, stable IC indicates that the model is capturing real signal. Portfolio managers and quants use IC as the primary diagnostic for cross-sectional ML strategies.

Feature importance/SHAP values. Interpretability is critical for governance. Feature importance metrics (gain, cover, frequency) and SHAP values attribute predictions to signals like momentum, carry, or skewness. Risk teams and CIOs rely on these diagnostics to ensure models are not "black boxes."

Cross-horizon correlation. By measuring the correlations between short-, medium-, and long-horizon models, practitioners can quantify diversification benefits. Risk managers monitor these matrices to validate ensemble robustness.

Factor alignment score. Machine learning outputs should align with known commodity factors. By measuring the correlation between ML signals and canonical factors, practitioners ensure economic consistency. Strategists and CIOs use this metric to validate that ML strategies are grounded in theory.

Glossary

Ensemble modeling: Combining multiple models or prediction horizons into a single aggregate score to improve robustness and reduce fragility.

Feature engineering: The design and transformation of input variables (e.g., momentum, skewness, basis) into structured signals usable by ML algorithms.

Gradient boosted trees (GBT): An ensemble ML algorithm that builds shallow decision trees sequentially, each correcting the errors of the last; widely used for return prediction.

Interpretability tools: Diagnostic techniques (feature importance plots, SHAP values, LIME) used to explain ML outputs to risk committees, regulators, and boards.

Pseudo-betas: Using feature importance scores as factor loadings to measure how a model's reliance on specific signals evolves over time.

Regularization: Techniques that constrain model complexity (e.g., L1/L2 penalties, limiting tree depth) to reduce overfitting in noisy financial data.

Related Content

Cao, Larry, ed. 2023. *Handbook of Artificial Intelligence and Big Data Applications in Investments*. CFA Institute Research Foundation. https://rpc.cfainstitute.org/research/foundation/2023/ai-and-big-data-in-investments-handbook.

CFA Institute Research and Policy Center. 2025. "The Automation Ahead: Exploring the Next Frontier of Investment Automation." https://rpc.cfainstitute.org/research/the-automation-ahead-content-series.

Fabozzi, Frank J., Andrew Chin, Igor Yelnik, and Jim Kyung-Soo Liew. 2025. "The Disappearing Edge: AI, Machine Learning, and the Future of the Discretionary Portfolio Manager." *Financial Analysts Journal* 81 (4). https://rpc.cfainstitute.org/research/financial-analysts-journal/2025/disappearing-edge-ai-machine-learning.

Wilson, Cheryll-Ann. 2025. "Creating Value from Big Data in the Investment Management Process: A Workflow Analysis." CFA Institute Research and Policy Center. https://rpc.cfainstitute.org/research/reports/2025/creating-value-from-big-data-in-the-investment-management-process.